arXiv:1410.6044v1 [cs.LO] 22 Oct 2014

AbPress: Flexing Partial-Order Reduction and Abstraction

Daniel Kroening, Subodh Sharma, and Björn Wachter

Department of Computer Science, University of Oxford, UK

Abstract. Partial-order reduction (POR) and lazy abstraction with interpolants are two complementary techniques that have been successfully employed to make model checking tools for concurrent programs effective.
In this work, we present *AbPress* – **A**bstraction-**b**ased **P**artial-order **R**eduction with **S**ource-**S**ets – an algorithm that fuses a recently proposed and powerful dynamic POR technique based on *source-sets* and lazy abstraction to obtain an efficient software model checker for multi-threaded programs. It trims the interleaving space by taking the abstraction and *source-sets* into account. We amplify the effectiveness of *AbPress* with a novel solution that summarizes the accesses to shared variables over a collection of interleavings.
We have implemented *AbPress* in a tool that analyzes concurrent programs using lazy abstraction, *viz.,* Impara. Our evaluation on the effectiveness of the presented approach has been encouraging. *AbPress* compares favorably to existing state-of-the-art tools in the landscape.

1 Introduction

The generation of safety proofs for concurrent programs remains a major challenge. While there exist software verification tools based on abstraction that scale to sequential systems code [1], that cannot be said for multi-threaded software. Abstraction-based verification of sequential programs works by annotating program locations in the control-flow graph with safety invariants. However, applying similar techniques to concurrent software is ineffective, as interleavings lead to an explosion of the control-flow graph. Therefore, along with abstraction of data, techniques are needed to effectively deal with interleaving explosion.

Partial-order reduction (POR) [2–4], a path-based exploration approach, is a technique that addresses the explosion in the interleaving space. The key notion in POR techniques is the *independence* of actions. Independent actions can commute, resulting in interleavings that cause no observable change in the output. All interleavings obtained by commuting independent actions fall into the same equivalence class. Thus, exploring only representative interleavings results in a reduction in the number of interleavings explored in total. Consider the example in Figure 1. The first two steps of T_1 are independent with the first two steps of T_2 (since they write to different shared variables). Out of the two interleavings shown in Figure 2, POR will identify that the interleaving $\langle 00, 10, 11, 21\rangle$ does not need to be explored.

There are also path-based techniques to address the problem of data-state explosion. A prominent technique is lazy abstraction with interpolants (the Impact algorithm) [5].

Main	T_1	T_2

`x=0; y=0;`
`create(`T_1`);`
`create(`T_2`);`
`join(`T_1`);`
`join(`T_2`);`
`assert (`$y \geq 0$`);`
	`1: x =1;`
	`2: x=x+1;`
	`3: x=y;`
		`1: y=-1;`
		`2: y= y+1;`

Fig. 1: Example with races

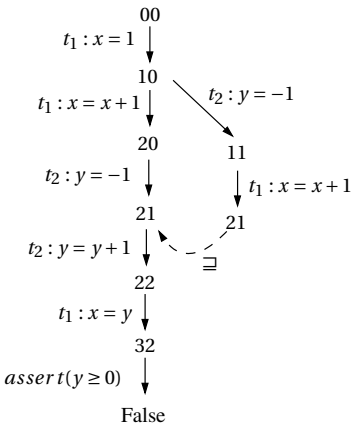

Fig. 2: Covers

The Impact approach begins by unwinding a program's control-flow graph into a tree. Each node in the tree (encoding the control location) is initially labeled with the state predicate `TRUE`, which indicates reachability of the node from the initial location. On reaching an error location, the node labels along the path to that node are updated with interpolants in order to prove that the error state is unreachable. The starting node is labeled `TRUE` and each subsequent node is assigned a formula that implies the next node's formula by executing the intervening program instruction. If the error node is labeled with `FALSE` then the approach has proved the path to be infeasible. The path exploration can terminate early. This happens when Impact discovers *covered* nodes. When two nodes v_1 and v_2 in the abstract reachability tree have the same program control location and the invariant at v_1 subsumes the invariant at v_2, then we say that v_1 *covers* v_2. This implies that it is no longer necessary to explore the reachability tree that follows v_2. For instance, in Figure 2 location 21 in the interleaving $\langle 00, 10, 11, 21\rangle$ has the same interpolant as the interpolant in location 21 of interleaving $\langle 00, 10, 20, 21, 22, 32\rangle$. Thus, the right 21 node is covered by the left 21 node. Observe that any implementation of POR would have eagerly detected the independence between $t_1 : x = x + 1$ and $t_2 : y = -1$ and the exploration of the right interleaving would have been avoided. The notion of covers is thus most useful when control-flow branching is present in the program.

Both Impact and POR, in particular dynamic POR (DPOR) [6, 7], use backtracking mechanisms to explore alternative choices at control locations: Impact uses backtracking for branching control flow and DPOR for interleavings. Due to the operational similarity and respective effectiveness in addressing problems arising from data and schedule explosion, Impact and DPOR are ideal candidates to be fused. Impara [8] offers a framework where Impact can be combined with a POR technique of choice. Impara comes with an implementation of the *Peephole POR* (PPOR) algorithm [9], which leaves room for further improvements. In particular, PPOR is known to be suboptimal for programs with more than two threads. Further, PPOR does not integrate a backtracking mechanism; it is a symbolic approach where chains of *dependent* actions have to be maintained at each node by ascertaining information from the future execution of the

program. DPOR algorithms can potentially be more efficient than PPOR. The recent work in [7] offers us an opportunity to use a dynamically constructed set of *dependent* actions, namely *source-sets*. Opportunities also exist to fine-tune the fusion of Impact and DPOR where the abstraction constructed by Impact feeds information into DPOR.

In this paper, we present a new verification algorithm for multi-threaded programs where Impact and DPOR with source-sets are combined in a novel way. Note that combining *covers* and DPOR in a sound manner is a non-trivial exercise. Consider Figure 3. Let n_C be the covering node and n_c be the covered node. Let n be the least common ancestor of nodes n_c and n_C.

To discover alternate schedule choices in p_1, DPOR will first enumerate the paths in the subtree from n_C. For each path p_s in the subtree, DPOR will perform a dependence analysis for each step in p_1 with each step in p_s (backtracking mechanism). Such an approach turns out to be prohibitively expensive. Therefore we *summarize* the accesses in the subtree and re-use the summary. This summarization technique is one key element to obtain an effective combination of DPOR and the covers Impact uses.

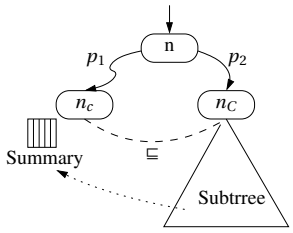

Fig. 3: Shared access summarization

Contributions: Our main contributions are: (1) an algorithm, *AbPress*, that combines source-set based DPOR with Impact, (2) abstract summaries of shared variable accesses in a subtree to create a sound fusion of DPOR and covers, and (3) a comparison of *AbPress* with the state-of-the-art tools in the landscape. We present the basic definitions associated with Impact in Section 2. We present the essentials of source-set DPOR in Section 3 and abstract summaries in Section 4. The complete algorithm *AbPress* is presented in Section 5. Experimental results are discussed in Section 6.

2 Preliminaries

We consider a concurrent program $\mathscr{P}$ that is composed of a finite set of threads $\mathscr{T}$. Each thread executes a sequence of operations given in C or C++. The threads communicate with each other by performing operations on shared communication objects such as global variables, semaphores and locks. We only consider programs with a fixed number of threads. A thread $T \in \mathscr{T}$ is a four-tuple $T = (L, l_0, A, l_\varepsilon)$ consisting of a finite set of program control locations L, an initial location $l_0 \in L$, a set of thread actions A and an error location $l_\varepsilon \in L$. A thread action a is a triple $a = (l, c, l')$ where $l, l' \in L$ are the entry and exit program locations for the action, respectively, and c is the program instruction. For brevity, we denote an action of thread T that is enabled at location l by $a_{T,l}$. We assume that we are working with an intermediate program representation where an instruction is either an assignment or an assume statement.

For notational convenience, we identify instructions using their standard formalisation as first-order formulae over the alphabet of primed and unprimed program variables $V \cup V'$. We denote the set of all such formulae by $\mathscr{F}(V \cup V')$. Consider the example in Fig. 1. For the assignment $z = 1$ in T_1, we have the action $(l_0, (z = 2 \wedge z' = z), l_1)$.

A *global control location* is a tuple with one component per thread, and is given as function $\bar{l} : \mathcal{T} \to L$. Let L_G be the set of all global control locations. By $\bar{l}[T \mapsto l]$, we denote the global location where the location of thread T maps to l while the locations of all the other threads remain unchanged. An action $a \in A$ from thread T is enabled if the action is enabled at $\bar{l}(T)$.

A program path π is a sequence $\pi = \sigma_0, \ldots, \sigma_N$ where $\sigma_i = (\bar{l}_i, T_i, a_i, \bar{l}_{i+1})$ consists of an action a_i from thread $T_i \in \mathcal{T}$ and a_i's entry and exit global program locations, $\bar{l}_i$ and $\bar{l}_{i+1}$. A path is an *error path* if $\bar{l}_0$ is initial control location for all threads, and $\bar{l}_{N+1}$ contains an error location of a thread.

We denote by $\mathscr{F}(\pi)$ the sequence of transition formulas $init^{(0)} \wedge R_0^{(0)}, \ldots R_N^{(N)}$ obtained by shifting each R_i i time frames into the future. Each R_i is a transition formula for an action at location $\bar{l}_i$. We say that π is feasible if $\bigwedge R_i^{(i)}$ is logically satisfiable. A solution for $\bigwedge R_i^{(i)}$ corresponds to a program execution and assigns values to the program variables at each execution step. The program is said to be *safe* if all error paths are infeasible.

2.1 Interpolants, Invariants and ARTs

In case a path is infeasible, an explanation can be extracted in the form of an interpolant. To this end, we recall the definition of *sequent interpolants* [10]. A sequent interpolant for formulas $A_1, \ldots, A_N$ is a sequence $\widehat{A}_1, \ldots, \widehat{A}_N$ where the first formula is equivalent to true $\widehat{A}_1 \equiv \mathit{True}$, the last formula is equivalent to false $\widehat{A}_N \equiv \mathit{False}$, consecutive formulas imply each other, i.e., for all $i \in \{1, \ldots, N\}$, $\widehat{A}_{i-1} \wedge A_i \Rightarrow \widehat{A}_i$, and the i-th sequent is a formula over the common symbols of its prefix and postfix, i.e., for all $i \in \{1, \ldots N\}$, $\widehat{A}_i \in \mathscr{F}(A_1, \ldots, A_i) \cap \mathscr{F}(A_{i+1}, \ldots, A_N)$. For certain theories, quantifier-free interpolants can be generated for inconsistent, quantifier-free sequences [10].

An *inductive invariant* is a mapping $I : L_G \to \mathscr{F}(V)$ such that $init \Rightarrow I(\bar{l}^i)$ (where $\bar{l}^i$ is the initial global control location) and for all locations $\bar{l} \in L_G$, all threads $T \in \mathcal{T}$, and actions $a = (l, R, l') \in T$ enabled in $\bar{l}$, we have $I(\bar{l}) \wedge R \Rightarrow I(\bar{l}[T \mapsto l'])$. A *safety invariant* is an inductive invariant with $I(\bar{l}) \equiv \mathit{False}$ for all error locations $\bar{l}$. If there is a safety invariant the program is safe.

Definition 1 (ART). *An* abstract reachability tree *(ART) $\mathscr{A}$ for program $\mathscr{P}$ is a tuple $(N, r, E, \sqsubseteq)$ consisting of a tree with nodes N, root node $r \in N$, edges $E \subseteq N \times \mathcal{T} \times \mathscr{F}(V \cup V') \times N$, and a covering relation $\sqsubseteq \subseteq N^2$ between tree nodes such that:*

- *every nodes $n \in N$ is labeled with a tuple $(\bar{l}, \phi)$ consisting of a current global control location $\bar{l}$, and a state formula ϕ. We write $\bar{l}(n)$ and $\phi(n)$ to denote the control location and annotation, respectively, of node n.*
- *edges correspond to program actions, and tree branching represents both branching in the control flow within a thread and thread interleaving. Formally, an edge is a tuple (v, T, R, w) where $v, w \in N$, $T \in \mathcal{T}$, and R the transition constraint of the corresponding action.*

We write $v \xrightarrow{T,R} w$ if there exists an edge $(v, T, R, w) \in E$. We denote $v \rightsquigarrow w$ if there is a path from v to w in $\mathscr{A}$. The role of the covering relation is crucial when proving program correctness for unbounded executions. It serves as an important criterion

in pruning the ART without missing error paths. The node labels, intuitively, represent inductive invariants that represent an over-approximation of a set of states. *Covering* relation, in other words, is the equivalent of a subset relation over this over-approximation between nodes. Suppose that two nodes v, w share the same control location, and $\phi(v)$ implies $\phi(w)$, *i.e.*, $v \sqsubseteq w$. If there was a feasible error path from v, there would be a feasible error path from w. Therefore, if we can find a safety invariant for w, we do not need to explore successors of v, as $\phi(v)$ is at least as strong as the already sufficient invariant $\phi(w)$. Therefore, if w is safe, all nodes in the subtree rooted in v are safe as well. A node is covered if and only if the node itself or any of its ancestors has a label implied by another node's label at the same control location.

To obtain a proof from an ART, the ART needs to fulfil certain conditions, summarized in the following definition:

Definition 2 (Safe ART). *Let $\mathscr{A} = (V, \epsilon, E, \sqsubseteq)$ be an ART. $\mathscr{A}$ is* well-labeled *if the labeling is inductive, i.e., $\forall (v, T, R, w) \in E :\ \bar{l}(v) = \bar{l}(w) \wedge \phi(v) \wedge R \Rightarrow \phi(w)'$ and compatible with covering, i.e., $(v, w) \in \sqsubseteq :\ \phi(v) \Rightarrow \phi(w)$ and w not covered. $\mathscr{A}$ is* complete *if all of its nodes are covered, or have an out-going edge for every action that is enabled at $\bar{l}$. $\mathscr{A}$ is* safe *if all error nodes are labeled with False.*

Theorem 1. *If there is a safe, complete, well-labeled ART of program $\mathscr{P}$, then $\mathscr{P}$ is safe.*

2.2 Path correspondence in ART

Let the set of program paths be Π_{CFG}. A program path $\pi \in \Pi_{CFG}$ is covered by $\mathscr{A}$ if there exists a corresponding sequence of nodes in the Π (denoting the set of paths in $\mathscr{A}$), where corresponding means that the nodes visits the same control locations and takes the same actions. In absence of covers, the matching between control paths and sequences of nodes is straightforward.

However, a path of the ART may end in a covered node. For example, consider the path $\langle 00, 10, 11, 21, 22 \rangle$ in the control-flow graph of Figure 2. While prefix $\langle 00, 10, 11, 21$ can be matched by node sequence $\langle v_{00} v_{10} u_{11} u_{21} \rangle$, node u_{21} is covered by node v_{21}, formally $u_{21} \sqsubseteq v_{21}$. We are stuck at node u_{21}, a leaf with no out-going edges. In order to match the remainder of the path, our solution is to allow the corresponding sequence to "climb up" the covering order $\sqsubseteq$ to a more abstract node, here we climb from u_{21} to v_{21}. Node v_{21} in turn must have a corresponding out-going edge, as it cannot be covered and its control location is also $\bar{l}_2$. Finally, the corresponding node sequence for $\langle 00, 10, 11, 21, 22 \rangle$ is $\langle v_{00} v_{10} u_{11} v_{21} v_{22} \rangle$.

This notion is formalized in the following definition:

Definition 3 (Corresponding paths & path cover). *Consider a program $\mathscr{P}$. Let $\mathscr{A}$ be an ART for $\mathscr{P}$ and let $\pi = (\bar{l}_0, a_{0,\bar{l}(0)}, \bar{l}_1) \ldots (\bar{l}_{N-1}, a_{m,\bar{l}(m)}, \bar{l}_N)$ be a program path. A* corresponding path *for π in $\mathscr{A}$ is a sequence $v_0, \ldots, v_n$ in $\mathscr{A}$ such that, for all $i \in \{0, \ldots, N-1\}$, $\bar{l}(v_i) = \bar{l}_i$, and*

$$\exists u_{i+1} \in N :\ v_i, \xrightarrow{T_i, R_i} u_{i+1}, a_i = (\bar{l}_i, R_i, \bar{l}_{i+1}) \wedge (u_{i+1} = v_{i+1} \vee u_{i+1} \sqsubseteq v_{i+1})$$

A program path π is covered by $\mathscr{A}$ if there exists a corresponding path $v_0, \ldots, v_n$ in $\mathscr{A}$.

Proposition 1. *Let $\mathscr{P}$ be a program. Let Π be a representative set of program paths. Assume that $\mathscr{A}$ is safe, well-labeled and covers every path $\pi \in \Pi$. Then program $\mathscr{P}$ is safe.*

We denote the set of enabled actions from a node $n \in N$ by $enabled(n)$. The edge from node n is denoted by $E(n)$. For any action a, let $proc(a) = T$ return the thread executing the action. We identify the unique successor node obtained after firing a from n by $a(n)$. In any given node $n \in N$, let $next(n, T) = a_{T,\bar{l}(T)}$ denote the unique next action to be executed from thread T after n. For a path $\pi \in \Pi$, the action fired from node $n \in \pi$ is a_n.

3 Partial Order Reduction with Source-sets

The basis for reduction using POR is the *independence relation* among concurrent actions. Intuitively, two concurrent actions are independent if executing then in any order leads to the same final state. Thus, a path π' obtained by commuting adjacent independent actions in π is same in behavior as π. The equivalence class representing all behaviorally similar interleavings is commonly known as a *Mazurkiewicz trace* [11]. In other words, Mazurkiewicz traces represent the partial order among the events of an execution path. It suffices to explore only representative execution (or one linearization) of each Mazurkiewicz trace. In context of this work, it means that exploring representative paths in $\mathscr{A}$ will suffice.

Definition 4 (Independent actions). *Let S represent the set of all execution states of the program. Two actions a_1 and a_2 are* independent*, denoted by $a_1 \parallel a_2$, iff the following conditions hold for all $s \in S$:*

- *Enabled: if a_1 is enabled in s then a_2 is enabled in $a_1(s)$ iff a_2 is enabled in s and*
- *Commute: $a_1(a_2(s)) = a_2(a_1(s))$*

The definition of independence is impractical to implement (as it requires a universally quantified check over the state-space). In practice, easily-checkable conditions can be provided to determine dependence of two actions (denoted by $\not\parallel$): for instance, two actions that are concurrent at location $\bar{l}$ that acquire the same lock or access the same shared variable (with one action performing a write) are dependent. In our setting, we consider actions that are enabled at a global location $\bar{l}$ to be independent when they *commute*.

POR algorithms operate by first computing a subset of *relevant* enabled actions from a node and explore only the computed subset from a scheduled node. Some of the popular techniques to compute this subset are *persistent-set* and *sleep-set* techniques [12]. Briefly, a set P of threads is persistent in a node if in any execution from the node, the first step that is dependent with the first step of some thread in P must be taken by *some* thread in P. Sleep-sets, on the other hand, maintain, at each state, information about past explorations and dependencies among transitions in the state in order to prune redundant explorations from that state. An elaborate exposition on these topics is beyond the scope of this paper. For a detailed discussion on these techniques, refer [12].

Dynamic POR (DPOR) techniques [6, 7, 13] compute the dependencies on the fly. This leads to the construction of more precise persistent-sets, thereby resulting in potentially smaller state-graphs for exploration. The central concept in most DPOR algorithms is that of a *race*. DPOR algorithms check whether actions in a path are racing and if found racing then the algorithm tries to execute the program with a different schedule to revert the race. We use $<_\pi$ to denote the total order among the nodes in the path $\pi \in \mathscr{A}$. Let $\rightarrow_\pi$ be the unique *happens-before* relation over the nodes in the path $\pi \in \mathscr{A}$ such that $\rightarrow_\pi \subseteq <_\pi$. Formally, consider $u, v \in N$; if $u \rightarrow_\pi v$ then $u <_\pi v$ and $a_u \not\parallel a_v$.

Definition 5 (Race). *Two actions a_u and a_v from nodes u and v in a path $\pi \in \Pi$ are in a* race*, denoted by $u \lessdot_\pi v$, if the following conditions hold true:(i) $u \rightarrow_\pi v$ and $proc(a_u) \neq proc(a_v)$ and (ii) there does not exist a node $w : u < w < v$ and $u \rightarrow_\pi w \rightarrow_\pi v$.*

DPOR was first introduced with *persistent-sets* [6]. However, recently in [7], an optimal strategy to perform DPOR was presented. Instead of using persistent-sets, the optimal DPOR relies on a new construct, namely *source-sets*. Succinctly, a source-set S at a state s is a set of threads that must be explored from s such that for each execution E from s there is some thread $p \in S$ such that the first step in E dependent with p is by p *itself*. Unlike persistent-sets where the first dependent step with p is taken by some thread in the set, in source-sets the first dependent step with thread p is taken by p itself. This subtle difference can lead to smaller exploration choices from a state. Source-sets are persistent-sets but all persistent-sets are not source-sets. DPOR based on source-sets has demonstrated considerable savings over basic DPOR with persistent-sets [7].

We provide a brief demonstration illustrating the differences between source-sets and persistent-sets using the example in Figure 4 (borrowed from [7]). Consider the path $r_1.r_2.q_1.q_2$ from the initial node and the persistent-set $\{p, q\}$. Note that r_2 is dependent with p but thread r is not in the persistent-set. By the preceding explanation of persistent-sets, it implies the persistent-set at the initial node must also include r. Consider again the prefix $r_1.r_2.q_1.q_2$ from the initial node. Let the source-set be $S := \{p, q\}$. The first step in the prefix that is a dependent action with p is r_2; however, note that r_2 is mutually independent with the actions from the process q. Thus, by reordering, we obtain $r_1 r_2 q_1 q_2 = q_1 q_2 r_1 r_2$. According to the explanation of a source-set, it is now the case that the first step in the execution prefix dependent with a source-set entry q is take by q itself. Thus, the given source-set S is sufficient to explore all executions starting from the start state. By contrast, the persistent-set definition mandated that $p, q,$ and r are explored from the start state.

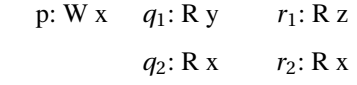

Fig. 4: Example for Source-sets

For a path $\pi \in \Pi$ starting from node n, let $I_n(\pi)$ denote a set of threads that have no *happens-before* predecessors in π. Intuitively, these are the "first steps" from threads $p \in I_n(\pi)$ at nodes $u \in \pi$. That is, there exists no $v \in \pi, v \neq u$ and $v \rightarrow_\pi u$. Let $WI_n(\pi)$ be the union of $I_n(\pi)$ and the set of processes $p \in enabled(n)$ such for all actions a in π, we have $next(n, p) \parallel a$. The set of threads $WI_n(\pi)$ represents the threads that can independently start an execution from the node n covering all possible paths from n.

Definition 6 (Source-sets). *A set* $\text{SSET}(n)$ *is a source-set for the set of paths* Π *after node* n *if for each* $p \in \Pi$ *we have* $WI_n(p) \cap \text{SSET}(n) \neq \emptyset$.

Our source-set based algorithm is similar to Algorithm 1 in [7]. However, unlike the version in [7], our version of source-set DPOR operates in a symbolic execution engine. Procedure $\text{COMPUTEBT}(u, v)$ in Algorithm 1 calculates the source-set SSET at node u when $a_u \lesssim_\pi a_v$ incrementally. Procedure NOTDEP(U,V) is the sequence of nodes π from the path $u \rightsquigarrow v$ (excluding u and v) such that each node w in the sequence is independent with u, *i.e.*, $u \nrightarrow_\pi w$.

4 Summarization

Combining source-set DPOR and lazy abstraction in a naive manner can lead to unsoundness. Consider Figure 3. Impact with DPOR will explore p_1, compute the relevant backtrack choices for the steps within p_2, and finally stops exploring any further since $n_c \sqsubseteq n_C$. However, a subset of paths following n_C will also follow from the node n_c. Terminating the dependency analysis without considering the dependencies among the shared variable accesses made in the sub-tree following n_C will result in relevant backtrack points in p_2 to be skipped. This is the source of unsoundness.

In order to be sound, the DPOR algorithm must be invoked for each path suffix in the sub-tree that follows a covering node n_C with each step in the prefix of the covered node n_c. Note that such a check quickly becomes expensive. We present an optimization of the above check by caching, for each shared variable, the set of threads that perform the "earliest" access to them.

From before, an edge $e = (u, T, a, w)$ shifts the control from node u to node w on action a. Let signature of a node $Sig(e) = (t, R, W)$ be a tuple consisting of the owner thread, the set of shared variables that is read by a and the set of shared variables written by a. Let Π be the set of paths starting from node n to the final node, i.e., for any path of the form $n \rightsquigarrow w$ where w is the final node with no actions enabled.

Definition 7 (Path Summary). *Let* $\text{SUM}(p)$ *be the signature of path* $p = e.p'$ *with* $Sig(e) = (t, R, W)$ *where the following conditions hold:*

- *if* p' *is empty then* $\text{SUM}(p) = \{Sig(e)\}$
- *if exist* $(t', R'W') \in \text{SUM}(p')$ *such that* $t = t'$, *then* $\text{SUM}(p) = \text{SUM}(p') \setminus \{(t', R', W')\} \cup \{t, R \cup R', W \cup W'\}$
- *if exist* $(t', Rd'Wr') \in \text{SUM}(p')$ *such that* $t \neq t'$ *and* $R \cap R' \neq \emptyset$ *or* $W \cap W' \neq \emptyset$, *then* $\text{SUM}(p) = \text{SUM}(p')[(t', R', W') \mapsto (t', R' \setminus R, W' \setminus W)] \cup \{Sig(e)\}$

Definition 8 (Node Summary). *The summary of a node* $n \in \mathscr{A}$ *is defined as the set* $\mathbb{S}(n) = \bigcup_{p \in \Pi} \text{SUM}(p)$ *where* Π *is the set of paths that start with root node* n.

Theorem 2 (Soundness of Shared Access Summarization). *Let* $\pi_1 = u_1 \ldots u_n$ *and* $\pi_2 = v_1 \ldots v_m$ *be two paths such that* $u_n \sqsubseteq v_1$. *For each node* $u_i \in \pi_1$, $\text{SSET}(u_i)$ *computed with* $\mathbb{S}(v_1)$ *over-approximates* $\text{SSET}(u_i)$ *when computed for the path* $\pi_1.\pi_2$.

Proof. Assume that there exists a thread $t \in \text{SSET}(u_i)$ when computed on $\pi_1.\pi_2$ which is not present in $\text{SSET}'(u_i)$ when computed with $\mathbb{S}(v_1)$. Let the assumed entry be (T, a) from node v. Since v must race with node u_i, clearly a_v (must be the "earliest" action accessing the shared variables in a racing manner after a_{u_i} (from Definition 5). From the invariant of the constructive definition of $\text{SUM}(\pi_2)$, a_v is a part of $\text{SUM}(\pi_2)$ and therefore a part of $\mathbb{S}(v_1)$. This contradicts our assumption.

We overload the operator for racing nodes; if $u \lessdot_\pi v$, then $a_u \lessdot_\pi a_v$ and $Sig(E(u)) \lessdot_\pi Sig(E(v))$. Consider Figure 5. Let nodes v, w, z fire actions that have the earliest accesses to variables x, y in path p_2 and p_3, as shown in the figure. The summary at n_C is $\mathbb{S}(n_C) = \{(t_2, \{y\}, \{x\}), (t_3, \{\}, \{x\})\}$. Observe that $u \lessdot_\pi w$ and $u \lessdot_\pi z$; therefore, we perform the source-set analysis for the path $u \ldots w$ and $u \ldots z$ by computing $\text{SSET}(u)$. Suppose we discover that $t_2 \in I_u(u \ldots w)$. We then add t_2 as an alternate schedule choice to $\text{SSET}(u)$. It is possible that t_2 at u is disabled since there was no earlier node that updated the value of y to one. This indicates that $\text{SSET}(u)$ can potentially be overapproximate when computed with summaries.

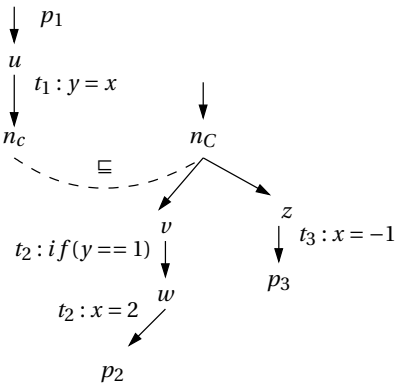

Fig. 5: Source-sets with summaries

5 AbPress Algorithm

ABPRESS is a combination of source-set DPOR with abstract summaries and Impact. We give the pseudo-code in Algorithm 1. A large part of Algorithm 1 is similar to Impara [8]. Functions BACKTRACK(V), CHOOSE(v) and COMPUTEBT(u, v) are the contributions of this work. We now give an overview of the algorithm.

A work list Q of nodes that are not fully explored is maintained along with the covering relation. Initially, Q contains the root node r and the cover relation is empty. EXPAND takes an uncovered leaf node and computes its successors. CHOOSE returns a thread that is chosen to be expored from a leaf node. We do not provide the algorithm for CHOOSE but briefly summarize its functionality. If the set of expanded threads and source-set from the node are empty, then any enabled thread is chosen, otherwise a thread from source-set is chosen. For every enabled action, it creates a fresh tree node w, and sets its location to the control successor l' given by the action. To ensure that the labeling is inductive, the formula $\phi(w)$ is set to *True.* Then the new node is added to the work list Q. Finally, a tree edge is added (Line 18), which records the step from v to w and the transition formula R. Note that if w is an error location, the labeling is not safe; in which case, we need to refine the labeling, invoking operation REFINE.

REFINE takes an error node v and, detects if the error path is feasible and, if not, restores a safe tree labeling. First, it determines if the unique path π from the initial node to v is feasible by checking satisfiability of $\mathscr{F}(\pi)$. If $\mathscr{F}(\pi)$ is satisfiable, the solution gives a counterexample in the form of a concrete error trace, showing that the program

Algorithm 1 AbPress

```
1: procedure MAIN()
2:     Q := {r}, ⊑ := ∅
3:     while Q ≠ ∅ do
4:         select and remove v from Q
5:         CLOSE(v)
6:         if v not covered then
7:             if error(v) then
8:                 REFINE(v)
9:             EXPAND(v)
10:    return 𝒫 is safe

11: procedure EXPAND-THREAD(T, v)
12:    (l̄, φ) := v
13:    for (l, a_{T,l̄(T)}, l') ∈ A(T) do
14:        w := fresh node
15:        l̄(w) := l̄[T ↦ l']
16:        φ(w) := True
17:        Q := Q ∪ {w}, N := N ∪ {w}
18:        E := E ∪ {(v, T, R, w)}

19: procedure BACKTRACK(v)
20:    π := r ... v path from r to v
21:    for u, w ∈ π : u <_π w do
22:        compute 𝕊(u)
23:        if Sig(u) ≪_π Sig(w) then
24:            COMPUTEBT(u, w)
25:        if exists z : v ⊑ z then
26:            for e ∈ 𝕊(z) do
27:                if Sig(u) ≪_π e then
28:                    COMPUTEBT(u, g(e))

29: procedure COMPUTEBT(u, v)
30:    π' = r ... u path from r to u
31:    π'' = NOTDEP(u, v).v
32:    if I_u(π'') ∩ SSET(u) = ∅ then
33:        ∃t ∈ I_u(π'') : SSET(u) ∪= {t}

23: procedure EXPAND(v)
24:    T := CHOOSE(v)
25:    if T = ⊥ then
26:        BACKTRACK(v)
27:    else
28:        EXPAND-THREAD(T, v)

29: procedure CLOSE(v)
30:    for w ∈ PRE(v) ∧ w uncovered : v ⊑ w do
31:        ⊑ := ⊑ ∪ {(v, w)}
32:        ⊑ := ⊑ \ {(x, y) ∈ ⊑ | v ⇝ y}
33:    if v covered then
34:        BACKTRACK(v)
35:

36: procedure REFINE(v)
37:    if φ(v) ≡ False then
38:        return
39:    π := v_0, ... v_N path from r to v
40:    if A_0 ... A_N = ITP(ℱ(π)) then
41:        for i = 0 ... N do
42:            φ := A_i^i
43:            Q := Q ∪ {w | w ⊑ v_i}
44:            ⊑ := ⊑ \ {(w, v_i) | w ⊑ v_i}
45:            φ(v_i) := φ(v_i) ∧ φ
46:        for w ∈ V s.t. w ⇝ v do
47:            CLOSE(w)
48:    else
49:        abort (program unsafe)
```

is unsafe. Otherwise, an interpolant is obtained, which is used to refine the labeling. Note that strengthening the labeling may destroy the well-labeledness of the ART. To recover it, pairs $w \sqsubseteq v_i$ for strengthened nodes v_i are deleted from the relation, and the node w is put into the work list again.

CLOSE takes a node v and checks if v can be added to the covering relation. As potential candidates for pairs $v \sqsubseteq w$, it only considers nodes created before v, denoted by the set $V^{\prec v} \subsetneq V$. This is to ensure stable behavior, as covering in arbitrary order may uncover other nodes, which may not terminate. Thus, only for uncovered nodes

$w \in Pre(v)$, it is checked if $\bar{l}(w) = \bar{l}(v)$ and $\phi(v)$ implies $\phi(w)$. If so, (v, w) is added to the covering relation $\sqsubseteq$. To restore well-labeling, all pairs (x, y) where y is a descendant of v, denoted by vE^*y, are removed from $\sqsubseteq$, as v and all its descendants are covered. Finally, if v is covered by z, BACKTRACK on v is invoked. The backtrack function performs the classic dependence analysis of DPOR. For each pair of nodes u, w where u, w in $r \dots tov$ and u races with w we compute the source-sets by calling the function COMPUTEBT (Lines 21-24). The functionality of COMPUTEBT is responsible for computing source-sets and is similar to Algorithm 1 in [7]. Since v is covered by z, the BACKTRACK function performs race analysis of each step u in $r \dots v$ with each entry e in the summary of z (Lines 29-31). If u and e race then the COMPUTEBT function is invoked again (with a ghost node for e) to compute the thread that should be added in the source-set.

MAIN first initializes the queue with the initial node ϵ, and the relation $\sqsubseteq$ with the empty set. It then runs the main loop of the algorithm until Q is empty, i.e., until the ART is complete, unless an error is found which exits the loop. In the main loop, a node is selected from Q. First, CLOSE is called to try and cover it. If the node is not covered and it is an error node, REFINE is called. Finally, the node is expanded, unless it was covered, and evicted from the work list.

6 Experiments

The purpose of our experiments is twofold: we would like to demonstrate the effect of the techniques proposed in the paper, and evaluate the competitiveness of our tool with comparable tools. To this end, we compare ABPRESS (IMPARA with Source-set DPOR and summaries) with three different tools:

- THREADER [14], a proof-generating software verifier for concurrent programs. It is one of the few other tools that produce correctness proofs for concurrent programs.
- FMCAD'13 [8]: IMPARA with peephole partial-order reduction [9], which serves as a baseline to evaluate the benefit of partial-order reduction.
- `CBMC` (version 4.9) [15], to compare with bounded model checking. Note that CBMC does not generate proofs for unbounded programs.

We evaluate on benchmarks of the Software Verification Competition [16] (SV-COMP 2014) and on weak-memory Litmus tests (submitted to SV-COMP 2015):

- `pthread`: This category contains basic concurrent data structures, and other lock-based algorithms. There are three challenging aspects to this category. (1) The queue examples and the stack example contain arrays. (2) The synthetic programs include the Fibonacci examples, which require a very high number of context switches to expose the bug. (3) Some examples contain more than 10 threads.
- `pthread-atomic`: This category contains mutual-exclusion algorithms and basic lock functionality, which is implemented by busy-waits. This creates challenging loop structures. Some loops are unbounded, i.e., there exists no unwinding limit, and some loops are nested.

– `pthread-ext`: This category is primarily designed to test the capability of tools that can deal with a parametric number of threads, which we have indicated with ∞. IMPARA does not terminate without a thread bound in this case. We ran IMPARA with a thread bound of 5, as this is the minimal number of threads it takes to expose all bugs. This is the only category in which IMPARA is incomplete, while tools that support parametric verification such as THREADER have an advantage.
– `Litmus`: These are small programs that are used to detect weakenings of sequential consistency. The benchmarks are `C` programs that have been instrumented to reflect weak-memory semantics [17] by adding buffers. The high degree of nondeterminism makes them challenging to analyse.

We ran our experiments on a 64-bit machine with a 3 GHz Xeon processor. Table 1 gives an overview of the results. For each benchmark, we give the number of lines (LOC) and the number of threads. For CBMC and THREADER, we give the running time and a tick mark if the benchmark was solved successfully. For ABPRESS and FMCAD13, we provide the running time, the number of nodes $|V|$ in the ART, and the time spent for solving SMT queries. The **effectiveness of summarisation** is tested by switching summarisation off, and, instead, enumerating the set of paths represented by the summaries. Our experiments confirm that summarisation dramatically reduces the cost of dependency analysis.

Without summarisation, we observe an order-of-magnitude increase of the number of paths explored in dependency analysis compared to summarisation. For example, in `qrcu_true` covering nodes have on average of around 14 postfixes. This means on average 14 paths would have to be analysed every time a cover between nodes is detected. As a result, successful cover checks become expensive. However, the IMPACT algorithm relies on covers being both efficient to check and to undo. In practice, this leads to timeouts, primarily, in programs with loops. For example, the analysis of `qrcu_true`, and `stack_true` timed out after 900 s. Overall, this naive algorithm is not competitive with FMCAD'13.

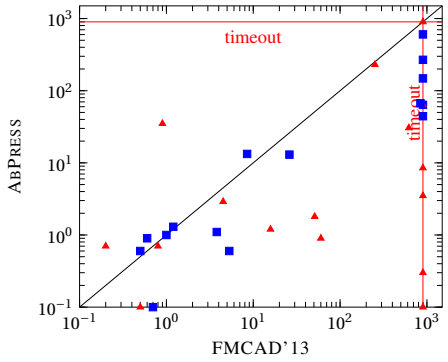

Fig. 6: AbPress vs. FMCAD'13

We evaluate the **benefits of Source-sets versus peephole partial-order reduction** by comparing against FMCAD'13. Figure 6 shows a scatterplot comparing the running times of FMCAD'13 with ABPRESS. The latter is clearly superior, resulting in both overall best running times and fewer timeouts.

As shown by Table 1, the number of ART nodes explored by ABPRESS is lower than for FMCAD'13, except in unsafe instances. As peephole POR explores more interleavings, it may by chance explore an interleaving with a bug earlier.

To evaluate the **competitiveness of AbPress**, as well as its limitations, we have aimed to carry out a comprehensive evaluation, where we deliberately retain examples

			CBMC	Threader	FMCAD'13			AbPress		
	LOC/Threads	safe	s	s	s	$\|V\|$	SMT	s	$\|V\|$	SMT
SV-COMP – pthread										
queue_ok_true	159/3	y	✓ 550.0	**ERR** –	**TO**	–	–	**63.7**	6489	14.7
queue_false	169/3	n	✓ 9.3	**ERR** –	**TO**	–	–	**8.5**	4867	1.5
stack_true	120/3	y	✓ 230.0	✓ 360.0	619	131507	336.5	**30.5**	7875	17.2
stack_false	120/3	n	✓ **0.5**	✓ 83.0	51	18776	28.8	1.8	2366	0.1
twostage_3_false	128/4	n	✓ 7.4	✓ 760.0	15.8	4290	12.5	**1.2**	3144	0.1
sync01_true	62/3	y	✓ 190.0	✓ **0.2**	**0.2**	731	0.0	0.7	775	0.3
sigma_false	48/17	n	✓ 30	**ERR** –	**TO**	–	–	**3.5**	692	0.1
indexer_true	83/14	y	✓ **1.4**	✓ 6.5	**TO**	–	–	**ERR**	–	–
reoder_2_false	84/3	n	✓ 1.4	✓ 2.4	60.1	11026	46.3	**0.9**	1005	0.2
reoder_5_false	2866/6	n	✓ **1.4**	✓ 2.6	**TO**	–	–	63.7	14546	42.3
lazy01_false	49/4	n	✓ 0.4	✓ 4.5	**TO**	–	–	**0.1**	147	0
bigshot_p_false	34/3	n	✓ 0.3	**ERR** –	0.5	272	0.0	**0.1**	147	0
bigshot_s_false	34/3	n	✓ **0.4**	**ERR** –	0.5	324	0.5	**WP**	144	0.1
bigshot_s_true	34/3	y	**WA** 0.4	**ERR** –	0.7	324	0.5	0.1	147	0.1
fib_bench_true	43/3	y	✓ 17.0	**ERR** –	**TO**	–	–	**TO**	–	–
fib_bench_false	40/3	n	✓ **1.0**	**ERR** –	**TO**	–	–	**TO**	–	–
SV-COMP – pthread-atomic										
scull_true	397/4	y	✓ **5.4**	✓ 610	**TO**	–	–	603.3	148,629	525
qrcu_true	147/3	y	✓ 850.0	**TO** –	**TO**	–	–	**268.8**	93742	219.7
qrcu_false	147/4	n	✓ **0.5**	**TO** –	0.9	1165	0.2	35.7	31453	19.1
dekker_true	54/3	y	✓ 120.0	✓ 3.2	1.0	883	0.7	**0.1**	331	0.0
peterson_true	41/3	y	✓ 2.7	✓ 5.3	0.6	746	0.4	**0.9**	1832	0.5
lamport_true	75/3	y	✓ 850.0	✓ 37	3.8	2560	3.1	**1.1**	3612	0.4
szymanski_true	54/3	y	✓ 7.4	✓ 13.0	**1.2**	1226	0.8	1.3	3098	0.7
read_write_lock_false	51/5	n	✓ **0.4**	✓ 22	5.3	4497	3.6	0.6	4899	53
read_write_lock_true	51/5	y	✓ **0.8**	✓ 17.0	842	93073	770.3	66.9	66041	16.9
time_var_mutex	54/3	y	✓ 2.4	✓ 2.6	**0.5**	1075	0.2	0.6	1196	0.2
SV-COMP – pthread-ext										
01_inc_true	47/∞	y	✓ 850.0	✓ **1.2**	26.1	433		13.0	148,629	3.8
02_inc_true	51/∞	y	✓ 850.0	✓ **3.9**	**TO**	–	–	44.3	93742	219.7
03_incdec_true	80/∞	y	✓ 850.0	✓ **13.0**	168.7	485808	47.3	123	31453	19.1
04_incdec_cas_true	99/∞	y	✓ 850.0	✓ **38.0**	**TO**	–	–	148.3	331	2
05_tas_true	57/∞	y	✓ 550.0	✓ 5.3	**TO**	–	–	**0.3**	1832	0.5
06_ticket_true	75/∞	y	✓ 850.0	✓ **0.8**	**TO**	–	–	**TO**	–	–
07_rand_true	97/∞	y	✓ 850.0	✓ 4.7	**TO**	–	–	**0.3**	3098	0.7
08_rand_case_true	123/∞	y	✓ 850.0	✓ 12.0	**TO**	–	–	**0.2**	3098	0.7
09_fmax_sym_true	59/∞	y	✓ 730.0	✓ **13.0**	**TO**	–	–	**TO**	–	–
10_fmax_sym_cas_true	69/∞	y	✓ 420.0	✓ **37.0**	**TO**	–	–	**TO**	–	–
Weak-memory and Litmus tests										
mix000_tso_false	359/3	n	✓ **1.1**	**ERR**	4.5	4126	2.5	2.9	3507	0.2
mix001_tso_false	519/3	n	✓ **2.75**	**ERR**	252.8	86812	209	23.0	15453	1.9
podwr000_power_opt_false	242/3	n	✓ 5.6	**ERR**	0.8	4740	0.2	**0.7**	2423	0.1
thin001_tso_true.c	194/3	y	✓ **1.1**	**ERR**	8.5	69961	2.9	13.3	30267	1.6

Table 1: Overview of benchmarks. The best time for each benchmark is in **bold font**. Results of IMPARA were obtained with SVN version 866. **WA** means that the tool produced a wrong alarm for a safe example. **WP** means that the tool produced a wrong proof for an unsafe example.

that are not main strengths of ABPRESS, e.g., where the number of threads is high or a very large number of thread interleavings is required to expose bugs.

ABPRESS **solves 10 SV-COMP benchmarks not solved by** THREADER. Two of those `qrcu_ok_safe` and `qrcu_false` are cases where THREADER times out. The other cases are errors where our implementation seems to be more robust in handling arrays and pointers[1]. Here the path-based nature of our algorithm can play out its strength in determining aliasing information. Furthermore, ABPRESS is capable of dealing with the weak-memory examples, where THREADER gives no results.

Disregarding the special `pthread-ext` category, the only cases where ABPRESS fails while THREADER succeeds are the pathological Fibonacci examples and the indexer example, which features 14 threads. Figure 7 compares the running times of THREADER with that of ABPRESS. The dots above the diagonal, where THREADER wins, are mainly in the `pthread_ext` category.

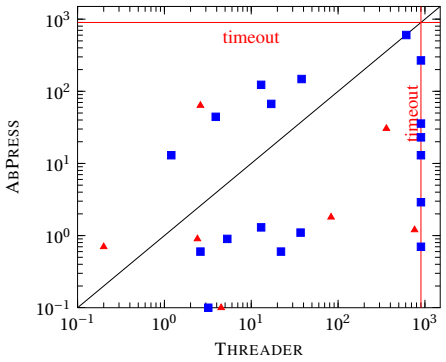

Fig. 7: AbPress vs. Threader

7 Related Work

Source-set based DPOR was recently presented in [7], as part of state-less explicit state model-checker for Erlang programs. While we borrow the notion of source-sets, our context is a fundamentally different. Hansen et al. consider a combination of partial-order reduction and zone abstractions for timed automata [18] where the dependence relation is computed from an abstract transformer.

Cimatti et al. [19] combine static POR with lazy abstraction to verify SystemC programs. Our work differs from their work on multiple fronts: SystemC has a significantly different concurrency model than multi-threading, and we use an abstract dynamic POR, which is inherently more precise than static POR.

We presented a combination of peephole partial-order with Impact in [8], however using peephole partial-order reduction which is simpler to integrate than source sets but leads to a greater number of interleavings, as demonstrated in our experiments.

THREADER is a software verifier for multi-threaded programs [14] based on compositional reasoning and invariant inference by constraint solving. In [20], Popeea et al present a combination of abstraction for multi-threaded programs with Lipton's reduction. Reduction is applied as a program transformation that inserts atomic section based on a lockset analysis. The authors then subsequently run THREADER on the transformed program. Unfortunately, at its current stage, their tool still requires manual transformations, and therefore we did not test against this implementation.

[1] ABPRESS gives only one wrong result, as it currently does not take failure of memory allocation into account, which affects example `bigshot_s_false`.

8 Conclusion

We have presented a concurrent program model checking technique *AbPress* that incorporates an aggressive DPOR based on source-sets along with Impact. Abstraction in the form of abstract summaries of shared accesses was utilized to amplify the effectiveness of DPOR with covers in the abstract reachability tree. We implemented the *AbPress* algorithm in Impara and evaluated it against comparable verifiers. Our initial results have been favorable. As a part of future work, we intend to use more aggressive property-guided abstractions to further reduce the interleaving space.

References

1. Ball, T., Levin, V., Rajamani, S.K.: A decade of software model checking with SLAM. Commun. ACM **54** (2011) 68–76
2. Peled, D.: All from one, one for all: on model checking using representatives. In: CAV. Volume 697 of LNCS. Springer (1993) 409–423
3. Godefroid, P., Wolper, P.: Using partial orders for the efficient verification of deadlock freedom and safety properties. In: CAV. (1991) 332–342
4. Valmari, A.: Stubborn sets for reduced state space generation. In: Advances in Petri Nets 1990. (1991) 491–515
5. McMillan, K.L.: Lazy abstraction with interpolants. In: CAV. (2006) 123–136
6. Flanagan, C., Godefroid, P.: Dynamic partial-order reduction for model checking software. In: Principles of Programming Languages (POPL), ACM (2005) 110–121
7. Abdulla, P., Aronis, S., Jonsson, B., Sagonas, K.: Optimal dynamic partial order reduction. In: Principles of Programming Languages (POPL), ACM (2014) 373–384
8. Wachter, B., Kroening, D., Ouaknine, J.: Verifying multi-threaded software with Impact. In: Formal Methods in Computer-Aided Design (FMCAD), IEEE (2013) 210–217
9. Wang, C., Yang, Z., Kahlon, V., Gupta, A.: Peephole partial order reduction. In: TACAS, Springer (2008) 382–396
10. McMillan, K.L.: An interpolating theorem prover. Theor. Comput. Sci. **345** (2005) 101–121
11. Mazurkiewicz, A.W.: Trace theory. In: Advances in Petri Nets. (1986) 279–324
12. Godefroid, P.: Partial-Order Methods for the Verification of Concurrent Systems – An Approach to the State-Explosion Problem. Volume 1032 of LNCS. Springer (1996)
13. Yang, Y., Chen, X., Gopalakrishnan, G., Kirby, R.: Efficient stateful dynamic partial order reduction. In: SPIN. Springer (2008) 288–305
14. Gupta, A., Popeea, C., Rybalchenko, A.: Threader: A constraint-based verifier for multi-threaded programs. In Gopalakrishnan, G., Qadeer, S., eds.: CAV. Volume 6806 of LNCS., Springer (2011) 412–417
15. Kroening, D.: CBMC. `http://cprover.org/cbmc` (2014)
16. Beyer, D.: Status report on software verification (competition summary SV-COMP 2014). In Abraham, E., Havelund, K., eds.: TACAS 2014. LNCS 8413 (2014)
17. Alglave, J., Kroening, D., Nimal, V., Tautschnig, M.: Software verification for weak memory via program transformation. In: ESOP. (2013) 512–532
18. Hansen, H., Lin, S., Liu, Y., Nguyen, T.K., Sun, J.: Diamonds are a girl's best friend: Partial order reduction for timed automata with abstractions. In: CAV. (2014) 391–406
19. Cimatti, A., Narasamdya, I., Roveri, M.: Boosting lazy abstraction for SystemC with partial order reduction. In: TACAS. Volume 6605 of LNCS., Springer (2011) 341–356
20. Popeea, C., Rybalchenko, A., Wilhelm, A.: Reduction for compositional verification of multi-threaded programs. In: Formal Methods in Computer-Aided Design (FMCAD). (2014) 187–194